



2171
\$

Patent

Attorney's Docket No. 033311-006

#7

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of)
)
J. Thomas NGO et al.) Group Art Unit: 2171
)
Application No.: 10/059,233) Examiner: Unassigned
)
Filed: January 31, 2002) Confirmation No. 6547
)
For: DATA REPLICATION BASED UPON A)
NON-DESTRUCTIVE DATA MODEL)

RECEIVED

FEB 25 2003

**PETITION TO MAKE SPECIAL
PURSUANT TO M.P.E.P. 708.02(VIII)**

Technology Center 2100

Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

In accordance with the provisions of M.P.E.P. 708.02(VIII), Applicants hereby request that the present application be granted special status and undergo accelerated examination. In support of this Petition, the following information is provided:

(A) A check in the amount of \$130.00 for the fee set forth in 37 CFR§1.17(h) is being submitted herewith;

(B) All claims are directed to a single inventive concept. Specifically, the application contains three independent claims, numbered 1, 109 and 114. These three claims are reproduced in Appendix A, which reflects the Preliminary Amendment being submitted herewith. All claims are directed to a method or system for replicating data at multiple devices. The independent claims each recite the core features of representing a data object as a graph of atoms at each device, where each graph comprises a parent atom with at least two descendant atoms, recording operations on the data object at a device by means of atoms in the graph, and updating the history of the data object at another device by sending the new atoms to the other device and/or forming a union of the graphs in the devices.

02/24/2003 SZEWDIE1 00000006 10059233

01 FC:1460

130.00 OP

(5/00)

(C) A online search was conducted for relevant patents using the keyword "replication". Patents dealing with biological replication were excluded by using the Boolean NOT operator in connection with the keywords "virus", "protein", "HIV", "cloning", "nucleic", "plasmid", "DNA", "repressor", "recombinant", "human", "viral", "retrovir" and "inhibitor". The most pertinent patents located in that search were provided in the Information Disclosure Statement filed April 3, 2002.

In addition, a Second Information Disclosure Statement is being submitted herewith, to provide information regarding commercial products that enable multiple users to collaborate with respect to a document.

(D) Copies of the patents cited in the Information Disclosure Statement filed April 3, 2002 are already of record. Copies of materials relating to the commercial products accompany the Second Information Disclosure Statement being submitted herewith.

(E) A detailed discussion of the references, pointing out how the claimed subject matter is patentable over the references, is presented herewith in Appendix B.

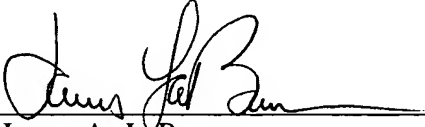
It is respectfully submitted that this Petition and the accompanying materials meet all of the requirements of M.P.E.P. 708.02(VIII). Granting the application special status is respectfully solicited.

Respectfully submitted,

BURNS, DOANE, SWECKER & MATHIS, L.L.P.

P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620

By: _____


James A. LaBarre
Registration No. 28,632

Date: February 20, 2003

for (703) 836-2021

Appendix A

Independent Claims

1. A method for replicating data at multiple devices, comprising the steps of:

- representing the history of a data object at each of said devices by means of a graph of atoms in a store, where said graph includes at least one parent atom of a first type that contains information pertaining to an operation performed on the data object and at least two descendant atoms of said first type;
- adding an atom of said first type to the atom graph in the store at a given device when an operation is performed on the data object at said given device; and
- updating the history of the data object at another device by transmitting to said other device at least one atom that is present in the store at said given device and absent from the store at said other device.

109. A method for replicating data at multiple devices, comprising the steps of:

- representing the history of a data object at each of first and second devices by means of a graph of atoms, where said atoms contain information pertaining to changes in the content of the data object, and said graph contains at least one parent atom having at least two descendant atoms;
- adding an atom to the atom graph at said first device when a change is made to the data object at said first device; and
- updating the history of the data object at a second device by forming a new graph at said second device that is a strict superset of the of the atom graph that existed at said second device prior to said updating step, and a non-strict subset of the union of the atom graphs that existed at said first and second devices prior to said updating step.

114. A system for replicating data at multiple devices, comprising:
a store at each of said devices, each of said stores containing replicas of a data object;
and
a manager associated with each store that detects and records operations performed on the data object in the form of a graph of atoms containing at least one parent atom having at least two descendant atoms, and that updates ^{complete change} the history of the data object by receiving atoms that are contained in a graph at one store and adding said atoms to the graph in another store.



APPENDIX B

Discussion of References

I. U.S. Patents

The 42 U.S. patents that were cited in the Information Disclosure Statement filed April 3, 2000 can be grouped into a few major categories that facilitate an understanding of the differences between their disclosures and the claimed invention.

RECEIVED

FEB 25 2003

A. Ancillary Aspects of Data Replication

Technology Center 2100

The pending claims are directed to the manner in which data objects, and updates thereto, are represented in a data-replication system, i.e. *how* data replication is implemented. The following patents generally relate to data replication systems, but are not concerned with the mechanism by which the data is actually replicated. Rather, they are directed to operations that are ancillary to the replication of the data itself.

6,253,211 - Gillies et al: Discloses the tracking of the replication of data and generation of alert messages upon a failure to replicate.

6,167,437 - Stevens et al: Discloses selective replication of data objects for efficient memory utilization.

6,167,427 - Rabinovich et al: Discloses selective replication of data for load-balancing purposes.

6,058,401 - Stamos et al: Discloses reducing the amount of data transmitted from one device to another during an update, to reduce communication overhead.

5,995,980 - Olson et al: Discloses the extraction of updates into batches for transmission.

5,991,768 - Sun et al: Discloses the suspension, or quiescence, of replication activities to perform administrative functions.

5,950,198 - Falls et al: Discloses a hashing technique for identifying identical files.

5,946,689 - Yanaka et al: Discloses the use of update serial numbers to detect contention.

5,859,790 - Sidwell: Discloses the replication of a data string in a register to generate a data string having a length that matches a predetermined bit capacity of the register.

5,838,921 - Speeter: Directed to the allocation of bandwidth resources in a client-server network.

5,832,225 - Hacherl et al: Directed to the management of replication topology information to facilitate adding servers to a network.

5,819,272 - Benson: Directed to the tracking of records in database replication.

5,787,442 - Hacherl et al: Pertains to directory services in a replication network.

5,771,355 - Kuzma: Directed to the transmission of e-mail over a network.

5,684,984 - Jones et al: Directed to the creation, management and publication of information to an online service.

5,265,164 - Matyas et al: Pertains to public key infrastructure.

B. Use of a Graph to Represent Revisions to an Object

1. Single Master Systems

One of the features of the claimed invention is the representation of operations on an object, e.g. changes, in the form of atoms in a graph. A particular advantage associated with this feature is the fact that the atoms can be replicated to all devices by a union operator, so that every device has a complete revision history. Thus, as successive revisions are made to data, more descendant atoms will be added to the graph, to represent those revisions. The following patents are directed to single master systems, wherein write operations can only occur at one device. In systems of this type, there is no need to keep the revision history at any of the other devices. As such, these systems do not have any need for a mechanism to combine revision histories across different devices, and hence do not suggest a concept such as representing objects by means of a graph of atoms having a parent atom with multiple descendants.

6,304,882 - Strellis et al

6,301,589 - Hirashima et al

6,167,438 - Yates et al

6,088,721 - Lin et al

5,960,460 - Marasco et al

5,884,324 - Cheng et al

5,745,753 - Mosher, Jr.

5,737,738 - Sharman

2. Multi-Site Revision Capability

The remaining patents disclose systems that support the ability to make changes to a data object at multiple sites, and to replicate those changes to the other sites. However, they do not disclose the use of a graph of atoms to record the changes and replicate the complete change history at all of the other sites. P,

6,182,117 - Christie et al: Discloses the use of e-mail messages to distribute data objects among sites. Each data object is identified by a unique identifier (UID), and if it is a revised version, the UID of the original object from which it evolved. This approach differs from the use of an atom graph, because the UIDs only represent the current and original versions of the object. They do not provide any of the intervening revision history, i.e. at most they constitute only a parent and one descendant.

6,138,124 - Beckhardt: Discloses the use of document sequence numbers and field sequence numbers to identify replicated data, to reduce bandwidth requirements. It does not describe atom graphs to represent the data sets at the nodes.

6,122,630 - Strickler et al: Discloses the transmission of database transactions from one node to another, and the use of a filtering mechanism so that only the transactions that occurred at the sending node are sent to the receiving node and those that originated at the receiving node are not re-sent. The transactions themselves are not representations of the data, rather they describe operations to be performed on the data. The actual data, i.e. the records in the database, are not represented by means of a graph containing at least one parent and two descendant atoms. *

6,049,809 - Raman et al: Discloses the use of cursors to identify the most current revisions received over multiple paths, to reduce bandwidth requirements. It does not disclose the use of atom graphs to represent data objects.

5,956,489 - San Andres et al: Discloses the use of an Arbiter to detect conflicting updates, and take offline any server whose update differs from that which the Arbiter determines should be the final outcome. It does not describe how data objects are represented, nor does it disclose how to record and combine revision histories.

5,937,414 - Souder et al: Discloses a hybrid system that utilizes synchronous updates for strong consistency and asynchronous updates for weak consistency. It does not

disclose how the data objects are represented, or whether a site stores complete revision history.

5,832,514 - Norin et al: Each node maintains a change set comprising a list of changes known at that node, and periodically broadcasts it to enable the other nodes to determine whether they are missing any changes, and if so request them. It does not disclose how a data object is represented at a node.

5,812,793 - Shakib et al: Discloses the use of change histories to identify the most recent version of an object. Only those changes which led up to the current version are identified in the change history, and thus it is different from a graph of atoms comprised of the changes that occur at multiple sites, whether or not those changes are reflected in the current version at a given node.

5,806,075 - Jain et al: Discloses transmission of both old and new values during replication for the purpose of detecting conflicts. It does not describe how the data is represented at the nodes.

5,799,306 - Sun et al: Discloses grouping of data objects as an alternative to schema level replication. It does not describe how the data objects are represented at the nodes.

5,794,253 - Norin et al: Discloses time-based expiration of data on a node-by-node basis. It does not describe atom graphs to represent the data sets at the nodes.

5,787,247 - Norin et al: Discloses the assignment of different states associated with the participation of a node in a data set, to prevent data loss. It does not describe atom graphs to represent the data sets at the nodes.

5,781,912 - Demers et al: Describes a mechanism for failure recovery by making changes permanent at a source and recording the changes as permanent at a destination once they have been made. If a failure occurs, the record is used to determine which permanent changes at the source site need to be sent to the destination site. It does not describe how the data is represented at the sites.

5,757,669 - Christie et al: Parent of 6,182,117, same disclosure.

5,737,601 - Jain et al: Related to 5,806,075, same disclosure.

5,588,147 - Neeman et al: Discloses the ability to replicate multiple data objects in a subtree, as well as single data objects. It does not describe how the data objects are represented.

5,261,094 - Everson et al: Describes a "pull" mechanism for replicating updates. It does not describe how the data is represented at the nodes.

4,620,276 - Daniell et al: Describes data objects in terms of atoms. It does not disclose that a data object is represented in the form of a graph of atoms.

II. Commercial Products

In addition to the foregoing patents, Applicants are aware of three commercial products that should be considered, since they provide replication of data.

A. Lotus Notes®

This product employs a hub-and-spoke system to synchronize data. A central server functions as the hub, and individual users form the spokes. At a given point in time, the hub and spokes may all contain the same version of the data, e.g. Version 1. Subsequently, User A may revise the data, which becomes Version 2 on his device. When User A connects to the hub, Version 2 is uploaded to the hub. Since Version 2 has a later time stamp, it supercedes Version 1. This version is then downloaded to each of the other spokes as they connect to the central server.

However, if User B also changes Version 1 before connecting to the hub, his device would contain yet a different version, e.g. Version 3. When he subsequently connects to the hub, a conflict occurs because both Version 2 and Version 3 supercede Version 1, but neither supercedes the other. To handle this situation, Lotus Notes creates a second file, which is replicated to all users. Thus, after synchronization, each user's device contains both Version 2 and Version 3 as separate data objects. In contrast to the claimed invention, therefore, Lotus Notes does not form a union of two graphs to represent a data object, so that a graph contains a parent atom with two descendants. Rather, it creates a second graph, with each graph comprising a parent having one descendant.

B. Groove®

This product is functionally similar to Lotus Notes, with the main difference being that it does not require a central server. Rather, any two devices can synchronize directly with one another. Like Lotus Notes, however, if Version 2 and Version 3 are independently generated from Version 1, the system handles the conflict by creating two separate data objects, rather than forming a union of the two versions. The two data objects are disseminated to all users as they connect to one another to synchronize. Thus, like Lotus Notes, a graph containing a parent with at least two descendants does not exist.

C. CVS

In this product, an official version of the data exists at a designated location, e.g. a central server, and working copies of this version are replicated on the users' devices. For example, the official version may be Version 1. If User A modifies his copy of this version, it becomes Version 2. If no changes are made to the official version in the meantime, then it is replaced with Version 2 when User A synchronizes.

However, if the official version was changed before synchronization, e.g. it became Version 3, the user has the responsibility for resolving all conflicts. In this case, User A merges his changes into the new official Version 3, to create Version 4. As long as no further changes have been made to the official version, Version 4 can then replace Version 3 as the official version. However, whenever the official version is changed while a user is working offline, he has the responsibility for merging his changes into the new official version before his changes can become part of the official version.

Thus, unlike the present invention, the data object at each device does not contain a full revision history, as results from a union of the various devices' atom graphs, where a graph can contain a parent with at least two descendant atoms. Rather, the individual devices contain at most two atoms, namely parent and child, to represent the data object.